LD-EnSF: Synergizing Latent Dynamics with Ensemble Score Filters for Fast Data Assimilation with Sparse Observations

Pengpeng Xiao[1,2] Phillip Si[1] Peng Chen[1*]

[1]Georgia Institute of Technology [2]Yale University

Code

ICLR

Introduction

Data assimilation (DA) integrates a **dynamical model** with **real-world observations**. Originally developed for weather forecasting, it is now widely used in geophysical and scientific modeling.

Why not rely on the model alone? Numerical models are imperfect for complex systems. Small initialization errors can grow rapidly, leading to significant deviations from the true state.

Why not rely on observations alone? Real-world measurements are noisy due to instrument limitations and sensor inaccuracies. In addition, observations are often sparse, and directly inserting them into the model can cause instability.

What DA does: DA systematically combines model predictions (prior) with observations to estimate a corrected posterior state. Applied iteratively, it accumulates information from past observations and continuously refines the system state over time.

Formally, let $x_t \in \mathbb{R}^{d_x}$ evolve as, with sparse and noisy observations ($d_y \ll d_x$):

$$x_t = F(x_{t-1}, u_t), \qquad y_t = H(x_t) + \gamma_t.$$

Goal: compute $P(x_t \mid y_{1:t})$ for all t. In the Bayesian filtering framework, this evolves $P(x_{t-1}|y_{1:t-1})$ to $P(x_t|y_{1:t})$ via two steps:

$$\textbf{Prediction:} \quad P(x_t|y_{1:t-1}) = \int P(x_t|x_{t-1})\, P(x_{t-1}|y_{1:t-1})\, dx_{t-1},$$

$$\textbf{Update:} \quad P(x_t|y_{1:t}) = \frac{1}{Z} P(y_t|x_t)\, P(x_t|y_{1:t-1}).$$

Iterating these two steps yields the full posterior trajectory.

Background

Classical DA methods such as EnKF and LETKF are efficient but assume linear, Gaussian posteriors, a poor fit for chaotic, nonlinear systems. Variational methods like 4D-Var are more accurate but require expensive repeated forward simulations.

Ensemble Score Filter (EnSF) [1] can handle nonlinear, high-dimensional posteriors. It samples the posterior via a reverse-time SDE (pseudo-time $\tau: 1 \to 0$, Gaussian $\to$ posterior). The posterior score is decomposed as:

$$\nabla_x \log P(x_{t,\tau}|y_{1:t}) = \underbrace{\nabla_x \log P(x_{t,\tau}|y_{1:t-1})}_{\text{prior score}} + h(\tau) \underbrace{\nabla_x \log P(y_t|x_{t,\tau})}_{\text{likelihood score}}, \quad h(\tau) = 1 - \tau$$

- The **prior score** is approximated via Monte Carlo using the ensemble.
- The **likelihood score** is computed analytically as $\nabla_x \log \mathcal{N}(y_t;\, H(x_{t,\tau}), R)$.
- No neural network training required.
- Under sparse observations the likelihood score vanishes in unobserved dimensions, causing EnSF to fail. **Latent-EnSF** [2] addresses this by assimilating in a VAE latent space.

However, both EnSF and Latent-EnSF propagate dynamics in the **high-dimensional full state space**, making the prediction step computationally expensive. **Surrogate models** offer a complementary solution by learning reduced dynamics directly in a **low-dimensional latent space**, avoiding repeated calls to the numerical simulator. Especially, **LDNets** [3]

- Outperform traditional surrogate models by achieving higher accuracy while using significantly fewer trainable parameters.
- Condition on **system parameters** u_t. Inaccurate parameters lead to significant prediction errors.
- Data assimilation can estimate and correct these parameters

LD-EnSF synergizes both: surrogate latent dynamics (LDNets) for fast forward propagation, and score-based filtering (EnSF) for accurate assimilation and parameter correction in the latent space.

LD-EnSF Framework

Latent Dynamics Network (LDNet) learns surrogate dynamics in a low-dimensional latent space ($d_s \ll d_x$): dynamics network $\mathcal{F}_{\theta_1}$, reconstruction network $\mathcal{R}_{\theta_2}$

$$\dot{s}_{t-1} = \mathcal{F}_{\theta_1}(s_{t-1},\, u_t), \qquad s_t = s_{t-1} + \Delta t\, \dot{s}_{t-1}, \qquad \tilde{x}(t,\xi) = \mathcal{R}_{\theta_2}(s_t,\, \xi).$$

Trained end-to-end via

$$\mathcal{L}(\theta_1, \theta_2) = \frac{1}{NMn} \sum_{j,t,\xi} \|\tilde{x}_j(t,\xi) - x_j(t,\xi)\|^2,$$

followed by fine-tuning of $\mathcal{R}_{\theta_2}$ with latent dynamics frozen.

LSTM Observation Encoder maps historical sparse observations to latent variables $(\hat{s}_t,\, \hat{u}_t) = \mathcal{E}_{\theta_3}(y_{1:t})$. Trained via

$$\mathcal{L}(\theta_3) = \frac{1}{Nn} \sum_{j=1}^{N} \sum_{t} \left(\|\hat{s}_t^{(j)} - s_t^{(j)}\|^2 + \|\hat{u}_t^{(j)} - u_t^{(j)}\|^2 \right).$$

Supports arbitrary observation locations and enables joint state–parameter assimilation.

EnSF in Latent Space is performed for both the latent states and parameters (s_t, u_t) using latent observations $(\hat{s}_t, \hat{u}_t)$.

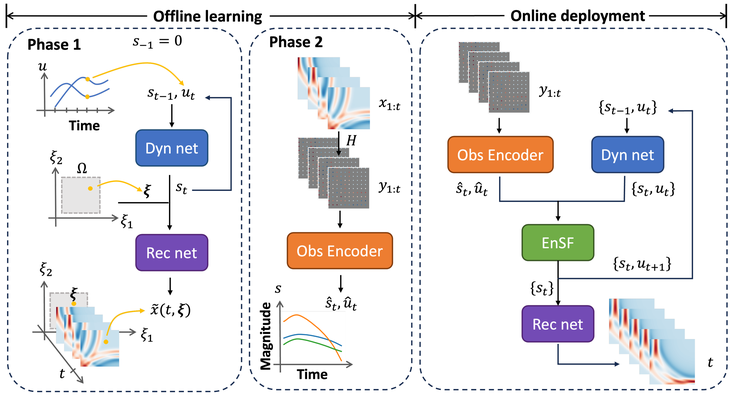

LDNet Surrogate Accuracy

Example	VAE	VAE-dyn	LDNet
Kolmogorov	0.0131	0.964	**0.0123**
Tsunami	0.0309	1.33	**0.0168**
Atmospheric	0.0856	0.483	**0.0656**

Table 1. Relative RMSE of surrogate model. LDNet achieves the lowest error in all cases.

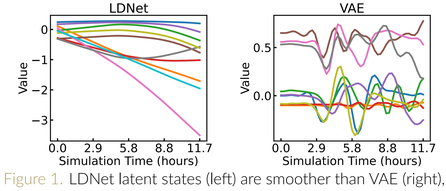

Figure 1. LDNet latent states (left) are smoother than VAE (right), enabling stable long-term dynamics.

Assimilation Results

Assimilation accuracy: With 0.44% observation density and 10% observation noise, LD-EnSF achieves lowest assimilation error.

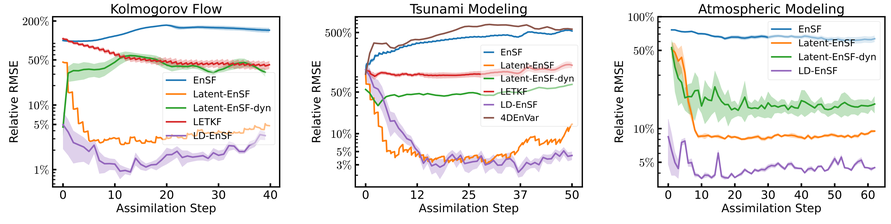

Computational speedup: LD-EnSF runs dynamics entirely in a compact latent space, eliminating full-space simulation cost.

	Kolmogorov Flow				Tsunami				Atmospheric			
Metric	LETKF	EnSF	L-EnSF	**LD-EnSF**	LETKF	EnSF	L-EnSF	**LD-EnSF**	LETKF	EnSF	L-EnSF	**LD-EnSF**
Dyn. time (s)	10829	10829	10829	**0.049**	211	211	211	**0.050**	35603	35603	35603	**0.070**
Filter time (s)	12729	40.13	0.71	**0.35**	10440	83.86	0.66	**0.37**	69820	27.36	2.42	**1.42**
Reconstr. (s)	–	–	6.86	**0.0018**	–	–	5.11	**0.0014**	–	–	35.08	**0.017**
Latent dim.	45K	45K	400	**10**	67.5K	67.5K	400	**12**	393K	393K	512	**52**

Table 2. Speedups of 2×10^5, 4×10^3, 5×10^5 on prediction step; assimilation also faster due to drastically reduced latent dimension.

Atmospheric visualization: Zonal wind velocity field from an atmospheric modeling with a given forcing field. Assimilation results and corresponding errors at the final time step.

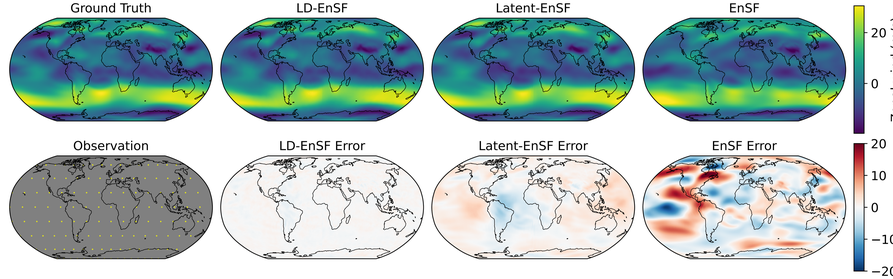

References

[1] Feng Bao, Zezhong Zhang, and Guannan Zhang. A score-based filter for nonlinear data assimilation. *Journal of Computational Physics*, page 113207, 2024.

[2] Phillip Si and Peng Chen. Latent-EnSF: A latent ensemble score filter for high-dimensional data assimilation with sparse observation data. *The Thirteenth International Conference on Learning Representations (ICLR)*, 2025.

[3] Francesco Regazzoni, Stefano Pagani, Matteo Salvador, Luca Dede', and Alfio Quarteroni. Learning the intrinsic dynamics of spatio-temporal processes through latent dynamics networks. *Nature Communications*, 15(1):1834, 2024.